

Autonomy Requirements for the Space Infrared Telescope Facility

G. Mark Brown¹, John C. Day²

¹Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Mail Stop 198-235
Pasadena, CA 91109
818-354-2215
g.mark.brown@jpl.nasa.gov

²Lockheed Martin Missiles and Space
1111 Lockheed Martin Way
B/107
Sunnyvale, CA 94089
408-742-9741
john.c.day@lmco.com

ABSTRACT

The Space Infrared Telescope Facility (SIRTF) will be launched in December 2001 into an Earth-trailing heliocentric orbit, receding from the Earth at a rate of approximately 0.1 AU per year. SIRTF carries three science instruments; taken together, these instruments provide imaging and spectroscopic capability at wavelengths from 3 μm to 180 μm . After an initial 60-day checkout and calibration period, SIRTF will spend the next five years conducting a variety of observing campaigns, including searches for brown dwarfs, superplanets, and planetary disks.

SIRTF's Ground System will schedule occasional one-hour communications sessions with the observatory, usually once or twice per day. In between communications sessions, SIRTF will operate autonomously, working its way through a ground-provided observing list that the ground will update once or twice per week. SIRTF's total science data return will be highly influenced by its ability to perform time-efficient observations. The helium supply has been sized to last five years, and the primary mission will end when that supply has been expended.

The limited frequency and duration of the planned flight-ground communications sessions, combined with the desire to perform time-efficient observations between those communications sessions, has motivated SIRTF to levy particular types of autonomy requirements. This paper will describe those autonomy requirements, and some initial plans for implementing those requirements.

1. SIRTF MISSION INTRODUCTION

The Space Infrared Telescope Facility (SIRTF), shown in Figure 1, is the fourth and last of NASA's Great Observatories. Scheduled for launch from Cape Canaveral in December 2001, SIRTF will combine the inherent sensitivity of a cryogenic space telescope with state of the art infrared detector technology. SIRTF's primary science objectives include:

- (1) the search for brown dwarfs and super-planets
- (2) the discovery and study of protoplanetary and planetary debris disks
- (3) the study of ultraluminous galaxies and active galactic nuclei
- (4) the study of the early and distant Universe




Figure 1. The SIRTF Observatory

A Delta II 7920H launch vehicle will deliver SIRTf to an Earth-trailing “solar orbit”, as shown in Figure 2. The observatory will move around the Sun in roughly the same orbit as the Earth, while drifting away from the Earth at a rate of roughly 0.1 AU/year. The solar orbit was chosen over a High Earth Orbit (HEO) primarily because it affords a larger launch mass; however the solar orbit also offers a more stable thermal environment, provides a more benign radiation environment, relaxes the viewing constraints, and doesn’t require a propulsion system for stationkeeping. For a good introduction to the SIRTf mission and its relationship to the other Great Observatories, see Reference 7.

2. OBSERVATORY DESIGN

The SIRTf flight system, henceforth referred to as the observatory, contains two engineering systems and three science instruments. The two engineering systems are the Spacecraft and the Cryogenic Telescope Assembly (CTA). The three science instruments are the Infrared Array Camera (IRAC), the Infrared Spectrograph (IRS), and the Multiband Imaging Photometer for SIRTf (MIPS).

The Spacecraft provides structural support for the CTA and the science instruments. The Spacecraft also provides the standard avionics functions for the entire observatory: power generation and distribution, command and data handling, pointing control, and telecommunications. The Spacecraft’s avionics subsystems, including its flight software, are largely inherited from the Mars Surveyor program. A significant exception is the pointing control subsystem, for which SIRTf requires greater pointing accuracy and stability than any of the Mars Surveyor spacecraft.

The CTA houses all of the telescope optics, including the 85 cm diameter primary mirror. The telescope is shielded from the Sun by the Spacecraft’s solar panel assembly, and is surrounded by an outer shell that radiates to cold space in the anti-Sun direction. The telescope is mounted to the forward dome of a vacuum shell. Inside the vacuum shell is a 360-liter superfluid helium tank and the Multiple Instrument Chamber (MIC), which houses all of the science instrument detectors. The helium tank provides a 1.4 K heat sink for the detectors, and the boiloff from this tank conductively cools the telescope to its operating temperature of approximately 5.5 K.

Inside the MIC, each of the three science instruments has one or more pickoff mirrors in the telescope’s focal plane. These mirrors redirect the incident infrared energy onto the appropriate detectors. The layout of the telescope’s focal plane, which has a




Figure 2: SIRTf’s Solar Orbit

diameter of 32 arcmin (a little more than 0.5 deg), is shown in Figure 3. There are ten different instrument apertures, all of which view the sky at all times; however only one of the instruments will be dissipating power and collecting data at any one time. All three instruments will rely on small, precise motions of the Spacecraft to move an observing target across a detector, or from one detector to another.

The Infrared Array Camera (IRAC) uses two 5.12 x 5.12 arcminute fields of view (FOVs). Each of the FOVs are split into two channels by a dichroic beamsplitter, with the 3.6 and 5.8 μm channels in one FOV and the 4.5 and 8.0 μm channels in the other. All four detector arrays have a 256 x 256 pixel format and each pixel has a physical size of 30 μm .

The Infrared Spectrograph (IRS) has four separate FOVs. Each FOV serves a dedicated detector, providing low resolution spectra ($R = \lambda/d\lambda = 50$) from 5 to 40 μm , and higher resolution spectra ($R = 600$) from 10 to 40 μm . A larger, fifth field of view, known as the peak-up array, is used to precisely locate sources so that they may be placed in the spectrograph slits within the required accuracy. This topic will be discussed in more detail in Section 5.

The Multiband Imaging Photometer for SIRTf (MIPS) uses five distinct optical trains to feed three detector arrays. Two of the arrays are used solely for photometry, one at 24 μm , and another at 160 μm . The third array provides photometry at 70 μm , as

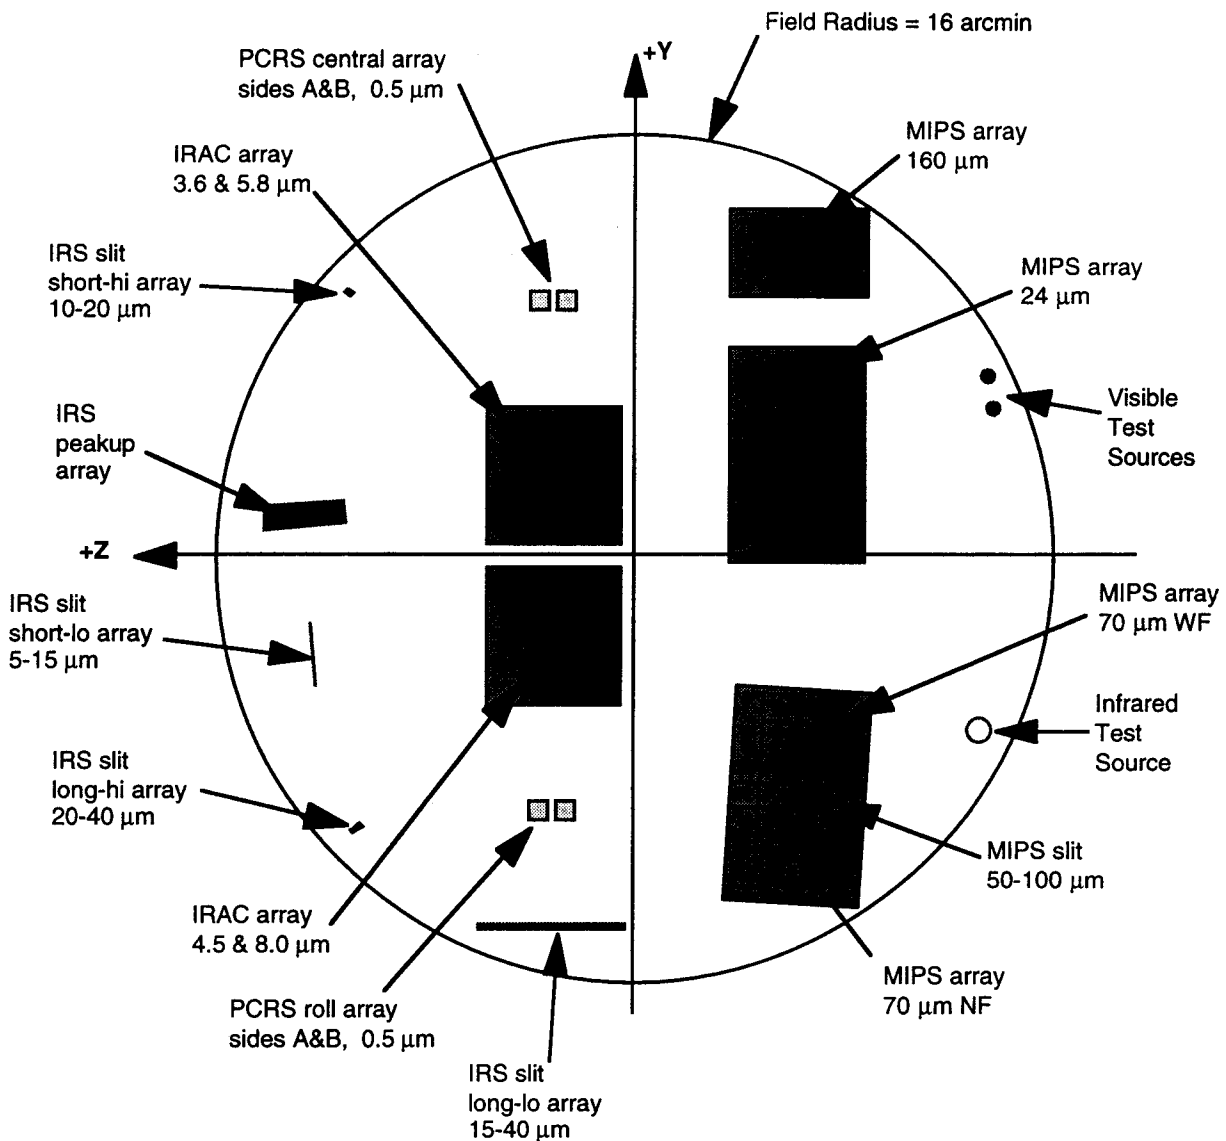


Figure 3. SIRTf Focal Plane Layout

well as very low resolution spectroscopy at 50 to 100 μm .

A redundant pair of Pointing Calibration Reference Sensors (PCRS) also have apertures in the focal plane. These engineering sensors are used by the Spacecraft's pointing control subsystem to estimate and compensate for misalignments between the telescope boresight and the star tracker boresight.

The observatory stands slightly less than 4 meters tall and has a maximum diameter of roughly 2 meters. The launch mass of approximately 900 kg includes 50 kg of liquid helium, which is expected to last at least five years.

3. OBSERVATORY OPERATIONS

The first sixty days of the mission are dedicated to In-Orbit Checkout (IOC) of the observatory. The Spacecraft is fully functional throughout this mission phase, but several activities are required to transition the CTA from its launch state to its fully operational state. These CTA activities include ejection of the dust cover, opening of the telescope's aperture door, and adjustment of the telescope's focus mechanism. Some of these CTA activities can only take place after the telescope has cooled down to its steady-state temperature, and the telescope isn't expected to reach that temperature until thirty days after launch. Flight-ground communications are almost continuous throughout IOC, and these CTA activities will be

directed from the ground using a combination of real-time commands and short, special-purpose sequences.

SIRTF's primary mission begins at the end of IOC, and is expected to last five years. During the primary mission, the observatory will be operated through ground-provided stored sequences that are typically uplinked once per week at a data rate of 2000 bps. Only one of the three instruments can be operated at any one time, but the stored sequence will direct the observatory to change instruments occasionally, perhaps as often as once per day.

Figure 4 illustrates SIRTF's viewing constraints. Pointing the telescope boresight within 80 degrees of the Sun causes direct solar illumination of the CTA, which cannot be tolerated. Pointing the telescope boresight more than 120 degrees from the Sun causes an unacceptable degradation of the solar array performance. The result is an annulus-shaped "operational pointing zone" that rotates as SIRTF revolves around the Sun. Targets near the ecliptic pole are always viewable; targets near the ecliptic

plane are viewable every six months for a period of at least 40 days.

The Spacecraft will regularly collect data from the active instrument, perform lossless compression of that data, and then store it in mass memory. Once or twice per day, the stored sequence will instruct the Spacecraft to stop collecting instrument data, point its fixed High Gain Antenna (HGA) at the Earth, and downlink its stored data to a Deep Space Network (DSN) antenna. The Spacecraft can store up to 4 Gbits in a single 12 hour observing period, which can then be downlinked in approximately 30 minutes to the DSN (at the maximum rate of 2.2 Mbps). The Spacecraft's mass memory is sized to accommodate one missed DSN pass without losing data.

SIRTF's primary mission will end when its superfluid helium supply has been expended. The telescope and the instruments will warm up to approximately 30 K, and the MIPS and IRS instruments will no longer be able to collect useful data. The IRAC's 3.6 and 4.5 μm bands could still be operated in an extended mission.




Figure 4. SIRTF's Viewing Constraints

4. THE MOTIVATION FOR AUTONOMY

Like many space missions, one of SIRTf's fundamental objectives is to maximize its Return on Investment (ROI). Several recent papers (e.g. References 3, 8) have championed autonomy as a means of increasing a mission's ROI, specifically by:

- (1) Reducing operations costs
- (2) Increasing the quality of its science data return
- (3) Increasing the quantity of its science data return
- (4) Ensuring the survival of the flight system in the presence of faults

Referring to the first ROI item, SIRTf has made little if any progress at employing autonomy to reduce operations costs. Some of the observatory's resources (e.g. the mass memory, the system's angular momentum) are autonomously monitored and/or controlled, and this alleviates a few minor ground tasks. However, the Ground System still has sole responsibility for operations planning, which is the largest operations expense.

The last three ROI items have been the dominant motivators for SIRTf's autonomous capabilities. Within a total SIRTf project cost of \$450 million, a 1% (\$4.5 million) expenditure on autonomy would be considered a "good buy" if it yielded any one of the following:

- A greater than 5% increase in the quality of science data returned
- A greater than 5% increase in the quantity of science data returned
- A greater than 5% increase in the reliability of the observatory

Having said all that, the SIRTf project has not created a specific autonomy budget, and has not justified its autonomous capabilities with quantitative cost/benefit assessments. However, we assert that the autonomous capabilities described in this paper each cost less than \$4.5 million, and each provides a greater than 5% benefit in one or more of the above-mentioned attributes.

This paper does not describe all of SIRTf's autonomy requirements, but the remaining sections provide the predominant examples of autonomy for each of the above-mentioned attributes. Section 5 describes the IRS peak-up process, an example of an autonomous function that enables higher quality science. Section 6 describes how autonomy is used to minimize the "dead time" between observations, increasing the total

number of observations that can be performed during the mission. Section 7 is an overview of the autonomous fault detection and recovery algorithms that permit the observatory to survive a failure.

5. TARGET ACQUISITION FOR THE INFRARED SPECTROGRAPH

SIRTf's Infrared Spectrograph (IRS) was described in Section 2. In order to perform high-resolution spectroscopy at infrared wavelengths, at least one dimension of the focal plane aperture needs to be extremely small. The short-wavelength aperture has the narrowest slit width, approximately 4 arcsec. As described in References 6 and 9, high-resolution spectrograph observations require the target to be positioned with an accuracy of 1/12th of the slit width slit, which is approximately 0.4 arcsec for the short-wavelength aperture.

Three uncertainties make it generally infeasible to directly position an observing target within an IRS slit with an accuracy of 0.4 arcsec:

- (1) The uncertainty in the Spacecraft's inertial attitude estimate, which is provided by the Spacecraft's star tracker, is approximately 5 arcsec.
- (2) The misalignment between the Spacecraft's star tracker and the telescope boresight will vary with time and temperature, and could change by as much as 0.4 arcsec in 24 hours.
- (3) The inertial coordinates of some of SIRTf's observing targets will not be known to an accuracy of 0.4 arcsec.

To overcome these uncertainties, a 60 arcsec x 60 arcsec FOV "peak-up" array has been added to the IRS. The Spacecraft will use real-time images from the peak-up array to first determine the target's position in the focal plane, and to then move the target onto the desired IRS slit. Figure 5 illustrates the steps of this autonomous process:

- (1) Through the stored sequence, the Ground System provides the Spacecraft with the inertial coordinates of a Local Position Reference (LPR), a companion "guide star" that has been selected because it radiates in the energy band at which the peak-up array is sensitive, and because its inertial coordinates are well known. The Ground System also provides the Spacecraft with the inertial offset between the LPR and the center of the observing target.

Observatory Autonomously Calculates Offset Slew $D = A + B + C$

where:

A is provided by the Ground System
 B is measured by the Peak-Up Array
 C is known a priori by the Observatory (requires in-flight calibration)




Figure 5. IRS Peak-Up Process

- (2) The Spacecraft plans and executes a slew that places the LPR inside the peak-up array's field of view.
- (3) The Spacecraft commands the IRS to take an image from the peak-up array, and to calculate the centroid coordinates of the brightest spot on the peak-up array.
- (4) The IRS returns the centroid coordinates to the Spacecraft, with a required accuracy of 0.15 arcsec.
- (5) Using the IRS-provided centroid coordinates, the ground-provided offset between the LPR and the observing target, and the known offset between the center of the peak-up array and the center of the IRS slit, the Spacecraft plans and executes an offset slew that places the observing target in the IRS slit.

Executing these offset slews with sub-arcsec accuracy, and then holding the slit stable on the target, are difficult pointing control challenges that are beyond the scope of this paper. For more information on SIRTf's pointing control subsystem requirements and design, see References 1, 2, and 6.

6. SEQUENCING

The presence of a consumable on which the bulk of the SIRTf science depends places an efficiency requirement on performing observations. The intent of

SIRTf's autonomous sequencing requirements is to perform the uplinked observations in the shortest time possible, without placing an undue burden on the observatory or removing appropriate functionality from the Ground System. With this in mind, the planning and execution responsibilities were split between the flight and Ground System. The Ground System retained the selection and prioritization of observations, and also the responsibility for determining the optimal ordering of these observations. The observatory was given the execution-time responsibilities: a simple activity selection function and the time compression function of event-driven execution. Within the observatory, requirements to implement these functions were assigned to both the sequence engine and design of the on-board sequences.

Event-driven execution (EDE) is a term which describes the behavior of an activity. SIRTf's activities are designed so that commands within the activity are spaced by relative-time timetags. Each command will execute a given number of seconds after the previous command was executed. However, EDE allows sequences to monitor variables provided by other portions of the flight software, blocking the executing sequence thread until the monitored variables meet specified criteria, indicating that an associated event has occurred. Typical uses of EDE are waiting for a slew to complete, or waiting for certain pointing stability criteria to be satisfied. EDE can also be used by one sequence to call another

sequence as a subroutine - stopping execution of the calling sequence until the called sequence completes.

EDE benefits the operation of the observatory in three ways:

- (1) It allows the Ground System to spend less time characterizing the time performance of observatory activities.
- (2) It allows the observatory to discontinue activities of low value (such as an observation in which the pickup operation fails to provide a centroid of adequate quality).
- (3) It allows the actual duration of an activity to be used by the on-board system in planning the next activity, rather than a ground-provided nominal duration plus a conservatively selected timing margin. The time saved by using the actual duration of an activity as opposed to the worst-case predicted duration of an activity can then be used to improve the observing efficiency of the observatory. The on-board activity planning function was specified to take advantage of this behavior.

SIRTF's DSN Contact activity, shown in Figure 6, is a good example of EDE. The DSN Contact activity is

an observatory-resident flight sequence that nominally establishes a two-way link with the DSN station, performs a playback of previously collected science data, and returns the observatory back to an observing mode. The sequence engine uses two real-time events during the execution of this activity:

- (1) First presence of an uplink signal from the DSN
- (2) A flight-software-generated flag that is set when science data playback has completed.

In the absence of an uplink signal from the DSN station, the observatory will, after a predetermined timeout period, resume science observations without transmitting its science data. This assumes that there has been some failure at the DSN station which prevents the contact, and therefore the observatory should return to science data collection. The on-board data storage is sized to be twice the expected amount of science data to accommodate this situation. However, collected science data is compressed using the lossless Rice algorithm, and the compression ratio is image-dependent. Since the amount of science data stored is variable, the actual time to playback the data may be less than the planned time. EDE permits the DSN Contact activity to recognize when the playback is complete, and to use this indication to stop the activity and return to science observation.

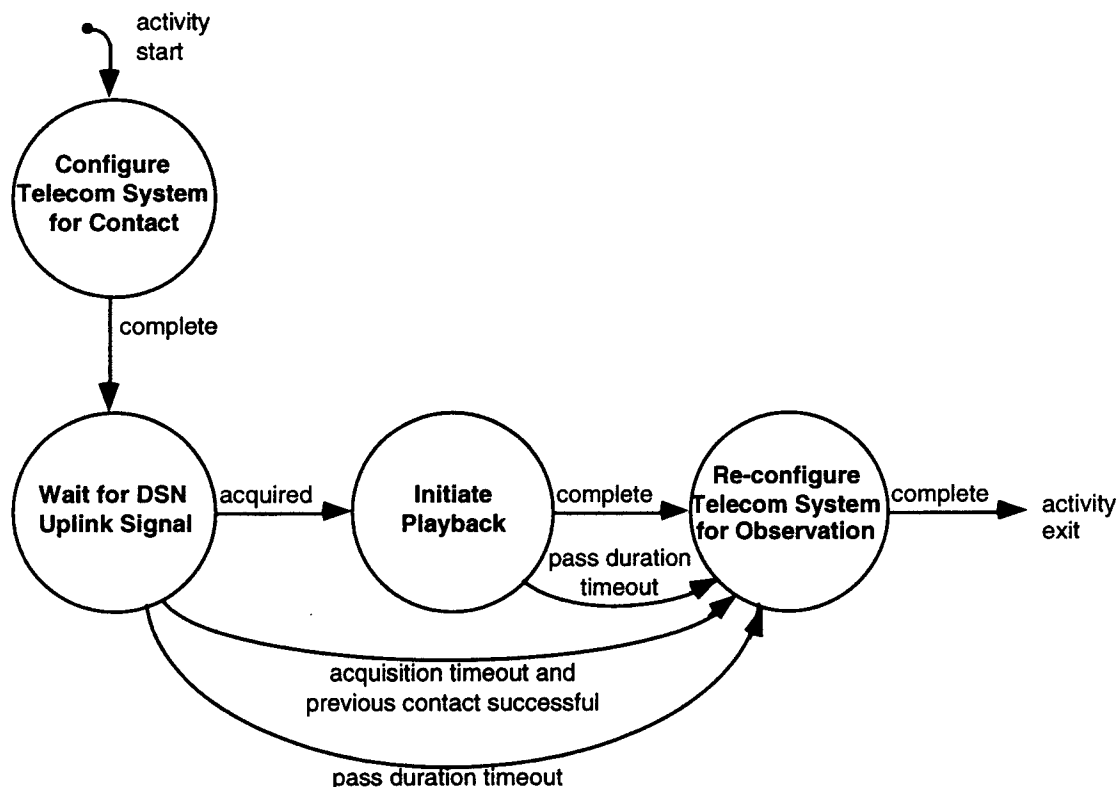


Figure 6. Deep Space Network (DSN) Contact Activity

The sequence engine's activity planning function is really an activity merging function. The Ground System uplinks to the observatory a set of absolute-time activities and a set of relative-time activities. The difference between an absolute-time activity and a relative-time activity is the desired start time of the activity. An absolute-time activity must start precisely at the time specified in the activity specification; a relative-time activity can start upon completion of any previous activity (whether absolute or relative). Both sets of activities are ordered by the ground before they are uplinked to the observatory. The observatory does not have the authority to alter the ground-specified order of the relative-time activities, or the ground-specified order of the absolute-time activities.

The Ground System produces meta-data associated with each activity to provide the observatory with the required information to merge the activities. The meta-data includes:

- (1) the desired attitude of the observatory at the start of the activity
- (2) the planned duration of the activity
- (3) the start time of each absolute-time activity

The relationship of this data to the elements of the sequence architecture is shown in Figure 7.

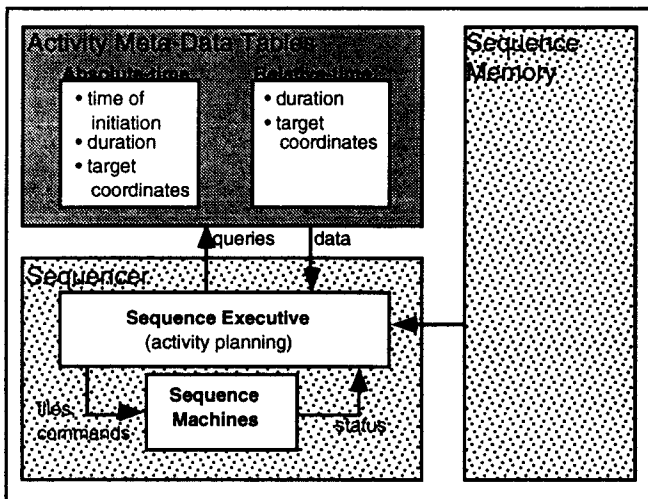


Figure 7. Sequence Architecture

The observatory uses this information to merge the relative-time activities with the absolute-time activities. The following simple algorithm, shown in Figure 8, describes the on-board logic that is used to initiate activities:

- (1) upon completion of the previous activity, determine if the next relative-time activity can be completed before the start time of the next absolute-time activity.
- (2) if time exists, start the next relative-time activity
- (3) if time does not exist, slew to the attitude required for the next absolute-time activity, and wait to start that activity

Note that when calculating whether a relative-time activity can complete before the next absolute-time activity, the observatory must estimate the time required to slew from its current attitude to the attitude of the relative-time activity, and the time required to slew from that attitude to the attitude of the next absolute-time activity. More efficient and complex schemes could have been used to maximize the available time, but the above scheme was chosen to retain activity ordering (predictability) and minimize the development and test cost of the planning algorithm.

The activity planning algorithm enables the observatory to take advantage of actual activity durations, rather than rely on worst-case duration estimates, as would be the case if each activity start time was defined by the Ground System. A consequence of this function is that the number of relative-time activities performed between two absolute-time activities is not strictly known in advance. Also, since the merging of the activities can't be easily predicted, the design of each activity must be such that there are no state or time dependencies between the absolute-time activities and the relative-time activities.

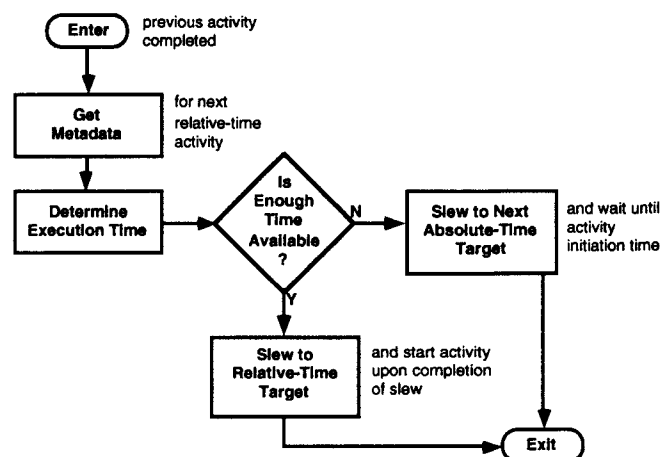


Figure 8. Activity Initiation Logic

7. FAULT DETECTION, ISOLATION, AND RECOVERY

SIRTF has two high-level project requirements that directly motivate autonomous fault detection, isolation, and recovery (FDIR). These two requirements are:

- (R1) "No credible single failure of the observatory shall result in the permanent:
- loss of data from more than one science instrument
 - loss of more than six months of its consumables, in particular the helium supply
 - loss of pointing control
 - loss of science-critical engineering telemetry
 - loss of capability to command the spacecraft"
- (R2) "The observatory shall have an efficiency of more than 90% in executing observations..... recovery from faults is considered inefficiency."

As described in Section 2, the Ground System will communicate with the observatory almost continuously during the 60-day In-Orbit Checkout period, and then will have 30-to-60 minute communications sessions once or twice per day during the primary mission. During certain observing campaigns, the time between communications sessions could be stretched to two or three days; nothing greater than that is envisioned.

Applying appropriate margin to this expected frequency of ground interaction, another key requirement was derived:

- (R3) "The observatory shall survive up to 7 days of unattended operation."

In other words, the observatory must perform autonomous FDIR whenever it is necessary to ensure compliance with (R1), assuming no help from the Ground System for up to 7 days.

The following subsections describe the resulting requirements for autonomous FDIR across the different flight systems of the observatory.

Autonomous FDIR for the Spacecraft

Many of the Spacecraft subsystems have employed block redundancy at the assembly level in order to enable compliance with (R1), subsequently referred to as the single fault tolerance requirement. For instance, SIRTF's Pointing Control Subsystem (PCS)

carries a redundant star tracker, a redundant Inertial Reference Unit (IRU), and an extra reaction wheel. Any of these redundant PCS assemblies can be used in combination with other prime PCS assemblies. SIRTF's Command and Data Handling (C&DH) Subsystem consists of two redundant VME chassis, each of which has cross-strapped data transport interfaces to the other engineering assemblies and the science instruments. SIRTF's approach to autonomous FDIR for these vital engineering subsystems is fairly similar to what's been done on previous large-scale deep space missions (see References 4, 5), and won't be described here in further detail. As on previous deep space missions, SIRTF will only exercise redundancy autonomously when it is necessary to survive a failure, where survival is defined by the single fault tolerance requirement.

From an operations perspective, the Spacecraft's autonomous fault responses fall into one of three categories. Figure 10 illustrates these response types, and the degraded modes of operation that are associated with them.

- A **Non-Interfering Response** permits all of the on-board stored sequences to continue. This type of response is always preferable, because no observing time is lost. However, this type of response can only be used in certain circumstances, when the response has no foreseeable interactions with the stored sequence, and will preserve all the system resources that were present at the time of the fault.
- A **Standby Response** halts all sequencing and initiates an immediate transition to **Standby Mode**. Upon entry to Standby Mode, the observatory terminates all active and future sequences, suspends the active science instrument (leaving its electronics powered), turns to a ground-prescribed communications attitude that is also thermally safe, and starts transmitting continuously through a ground-prescribed antenna. The rest of the observatory configuration remains unchanged, so that the state of the observatory at the next planned communications pass should be close to the expected configuration, simplifying the task of resuming operations. It is important to note that for almost the entire SIRTF mission, the observatory can remain thermally safe while pointing its HGA at the Earth, permitting the Ground System to utilize HGA communications in Standby Mode. In Standby Mode, the observatory maintains stellar reference, and can continue to point its HGA at the Earth using onboard ephemeris data.

- A **Safing Response** halts all sequencing and initiates an immediate transition to **Safe Mode**. Upon entry to Safe Mode, the observatory terminates all active and future sequences, points its +Z axis (the solar panel normal) directly at the Sun, powers off the science instruments and any other non-essential loads, and transitions to Low Gain Antenna (LGA) communications. In Safe Mode, the observatory keeps the +Z axis pointed at the Sun, and rotates slowly about the +Z axis. The Ground System should receive telemetry at least 50% of the time during each two-hour observatory rotation.

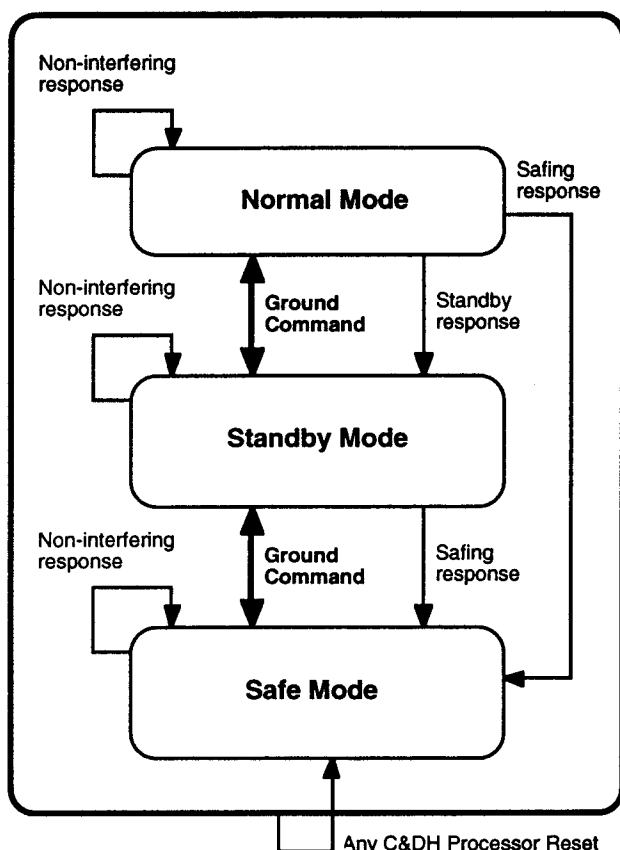


Figure 10: Autonomous Fault Responses

Figure 10 shows that prior execution of one type of fault response does not preclude future execution of a different type of fault response. However, once the observatory has entered a degraded mode, it cannot autonomously “improve” its mode. Ground System commands are required to resume science operations after Standby Mode entry or Safe Mode entry.

All autonomous FDIR activity by the Spacecraft will be time-tagged and logged in a “fault history table”. The fault history table will occupy a portion of the

“icebox” RAM that can be accessed by either of the flight computers. The fault history table will hold at least the first 50 activities and the last 50 activities that were entered. The contents of the fault history table will be downlinked regularly in the engineering telemetry stream.

Autonomous FDIR for the CTA

The CTA does not have its own flight processor; the Spacecraft’s flight processor controls the CTA’s mechanisms and monitors the CTA’s temperature and pressure sensors. The following subsections describe the autonomous FDIR requirements that have been assigned to the Spacecraft in order to protect the CTA from failures of either the CTA or the Spacecraft.

1. Assuring Successful Venting of the Helium Tank

The CTA’s helium tank cannot be vented until the atmospheric pressure is less than 1 Torr. In order for the tank to be vented safely, it should be vented while the tank is being accelerated by the launch vehicle, so that the helium is trapped away from the vent line. The time period during which both conditions are true is less than five minutes. SIRTf’s approach to this time-critical activity is as follows:

- (1) The CTA provides two redundant vent valves.
- (2) The Spacecraft receives three independent fairing separation indications from the launch vehicle.
- (3) The Spacecraft performs a two-out-of-three vote on the fairing separation indicators. Upon detecting fairing separation, the Spacecraft commands the primary vent valve to open via the primary CTA drive electronics, and commands the backup vent valve to open via the backup CTA drive electronics.
- (4) The Spacecraft verifies that at least one of the vent valves opened successfully, using limit switches on each of the valves, and the output of a single pressure transducer in the helium vent line. If the Spacecraft cannot verify that at least one of the vent valves opened, it will autonomously retry opening both valves.

2. Protecting the CTA from Direct Sun Exposure

The CTA is shielded from the Sun by the Spacecraft’s solar panel assembly, and is surrounded by an outer shell that radiates to cold space in the anti-Sun direction. The CTA design has been tailored for Sun-shaded operation; preliminary estimates show that if the Sun were to shine directly on its outer shell, the CTA would boil off one day’s worth of its helium supply for each minute of Sun exposure.

Since the observatory is three-axis stabilized, even a temporary loss of pointing control could cause an attitude excursion and undesired exposure of the CTA. After a thorough trade study, SIRTf has adopted the following requirements as its approach to tolerating temporary pointing control failures:

- (1) The Spacecraft will autonomously restrict the commanded attitude to a Sun-centered Operational Pointing Zone (OPZ). The approximate dimensions of the OPZ are shown in Figure 11. If the Ground System commands the Spacecraft to an attitude that is outside the OPZ, the Spacecraft will reject that command, remain at the last commanded attitude, and halt the affected observing sequence.
- (2) No single failure of the Spacecraft's pointing control subsystem will cause its attitude to deviate from its commanded attitude by more than 3 degrees about the X (roll) axis, and 5 degrees about the Y (pitch) axis. Figure 11 represents these possible excursions as a Pointing Recovery Zone (PRZ) that extends beyond the OPZ.
- (3) No single failure of the Spacecraft's command and data handling (C&DH) subsystem will disable pointing control for more than 20 seconds.
- (4) The Ground System will limit the commanded turn rates of the observatory to 0.10 deg/sec about the X (roll) axis, and 0.25 deg/sec about the Y (pitch) axis.
- (5) In order to accommodate unplanned attitude excursions from (2) or (3) in combination with the commanded rates from (4), the Spacecraft will shield the CTA from direct Sun exposure at all attitudes within the OPZ and the PRZ.

These requirements have not been arbitrarily assigned; they are a compromise solution, based on preliminary assessments of the current subsystem capabilities and the foreseen costs of enhancements. For instance, the Spacecraft will only be able to achieve (3) by using an enhanced, two-stage software initialization process. After a C&DH reset or side swap, the C&DH will need to immediately initialize the pointing control system and its software. Initialization of the rest of the software functions can be delayed until after the pointing control system has nulled the angular rates.

3. Protecting the CTA from High Power Dissipation

The superfluid helium supply has been sized to last for five years, assuming the science instruments dissipate an average of 8 mW in the Multiple Instrument Chamber (MIC). Each science instrument is allowed




Figure 11: Sun Safe Pointing Zones

to temporarily draw more power during selected activities, but autonomous corrective action is required if the power dissipated in the MIC exceeds a steady state value of $8 \text{ mW} \times (180 \text{ days} / 7 \text{ days}) = 200 \text{ mW}$. This is the rate at which the CTA would lose six months of its helium supply in less than one week. SIRTf's approach to protecting the CTA's helium supply is as follows:

- (1) The science instruments have primary responsibility to protect against unexpectedly high power dissipation. Each instrument is required to autonomously detect and isolate any failure that would cause it to dissipate more than 360 joules inside the MIC in a one hour period, which is an average of 100 mW during that period. Each instrument is expected to meet this requirement through some combination of localized current sensing and temperature sensing.
- (2) As a backup to (1), the Spacecraft will autonomously monitor a suite of CTA-provided temperature sensors inside the helium tank and inside the MIC. If a science instrument were to dissipate more than 100 mW steady-state inside

the MIC for any reason, including an operational error that left it in a high power state, some of these temperature measurements will rise slowly over long time scales (e.g. hours or days). When a majority of these measurements are consistent with abnormally high power dissipation, the Spacecraft will autonomously respond by removing all power from all science instruments.

4. Protecting the CTA from Low Power Dissipation

The average power dissipated in the MIC must remain above 2 mW in order to maintain unidirectional flow of the helium across the cryostat's porous plug. The Ground System has the primary responsibility to protect the CTA from against unexpectedly low average power dissipation. They are expected to model the power dissipated by each type of instrument activity, and to use the stored sequence to activate a 2 mW make-up heater if and when the commanded instrument operations are not expected to dissipate sufficient power. However, the Spacecraft has two related requirements:

- (1) As a backup to the Ground System, the Spacecraft is required to monitor the temperature differential across the CTA's porous plug; if this differential drops below a prescribed threshold, the Spacecraft will autonomously respond by turning on the above-mentioned 2 mW make-up heater.
- (2) If the Spacecraft autonomously removes power from the active science instrument as part of an autonomous fault response, the Spacecraft is also required to turn on the above-mentioned 2 mW make-up heater.

Autonomous FDIR for the Science Instruments

Two of SIRTf's science instruments, the MIPS and the IRS, share a common processor and other input/output electronics. In order to satisfy the single fault tolerance requirement, the MIPS/IRS common electronics are block redundant. The IRAC also has been permitted to fly block redundant electronics. All three science instruments will be self-monitoring to some extent, and will be able to report their errors to the Spacecraft. The Spacecraft will be responsible for verifying that the instruments are providing fresh data and are operating within acceptable temperatures and voltages. The Spacecraft will also be responsible for "safing" each instrument following either an instrument-reported error or a Spacecraft-detected error. Instrument recovery is not autonomous, because it is not justified by the single fault tolerance requirement; the Ground System will be responsible for instrument recovery, including the employment of redundant electronics.

8. SUMMARY

The autonomous capabilities of the Space Infrared Telescope Facility (SIRTf) will permit it to:

- (1) acquire observing targets for high-resolution spectroscopy
- (2) make the most efficient use of its consumables-limited observing lifetime
- (3) survive failures that could otherwise end its mission during periods of unattended operation

These autonomous capabilities have been baselined because of their predicted performance benefits, not because of any predicted savings in operations costs.

ACKNOWLEDGMENTS

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration; and at Lockheed Martin Missiles and Space, under contract with the Jet Propulsion Laboratory.

The authors would like to acknowledge the following individuals for their contributions to the development of SIRTf's operating strategies and sequencing architecture: Marie Deutsch, Mark Garcia, Ken Hooper, Joe Kahr, Paddy Lock, Mark Miller, David Mittman, Dennis Potts, and Steve Wissler. The authors would also like to thank Mark Garcia for providing Figures 1 through 4.

The authors would like to acknowledge the following individuals for their contributions to the development of SIRTf's fault protection architecture and requirements: Greg Andersen, Bill Clark, Marty Huisjen, Jeff Lee, Neil Martin, Mark Miller, Hank Mora, Dennis Potts, Rich Russek, John Troeltzsch, Nick Vadlamudi, and Steve Willner.

BIOGRAPHIES

G. Mark Brown is a Senior Engineer in the Avionics Systems Engineering Section at the Jet Propulsion Laboratory. He is currently the leader of SIRTf's Fault Protection Integrated Product Team, and is also the Flight Software Cognizant Engineer for Deep Space 3, a separated spacecraft optical interferometer that will be launched in 2003. Mark previously developed onboard fault protection algorithms for the attitude control systems of the Galileo spacecraft, which is currently orbiting Jupiter, and the Cassini spacecraft, which was launched to Saturn in October 1997. He holds B.S. and M.S. degrees in Aeronautics and Astronautics from the Massachusetts Institute of Technology, and a M.S. in Systems Management from the University of Southern California.

John C. Day is a Systems Engineer in the Remote Sensing and Space Science Business Unit at Lockheed Martin Missiles and Space. He is currently the flight operations lead at LMMS for the SIRTf Systems Engineering, Integration and Test contract. John previously worked as a system engineer at JPL on the Galileo and Cassini projects, primarily in information system and fault protection system requirements and verification. He received his B.S. in Electrical Engineering from Texas A&M University and his M.S. in Electrical Engineering (with an emphasis in Control Systems) from the University of California, Los Angeles.

REFERENCES

- [1] D. S. Bayard, "An Exposure-dependent Spacecraft Pointing Control Criteria for High-resolution Spectroscopy", 1998 American Control Conference, Philadelphia, PA, June 1998.
- [2] D.S. Bayard, T. Kia and J. Van Cleve, "Reconfigurable Pointing Control for High Resolution Space Spectroscopy," Proc. NASA University Research Centers (URC) Technical Conference on Education, Aeronautics, Space, Autonomy, Earth, and Environment, Albuquerque, NM, February 16-19, 1997.
- [3] D. E. Bernard et. al, "Design of the Remote Agent Experiment for Spacecraft Autonomy", IEEE Aerospace Applications Conference, Aspen, CO, March 1997.
- [4] G. M. Brown, D.E. Bernard, and R. D. Rasmussen, "Attitude and Articulation Control for the Cassini Spacecraft: A Fault Tolerance Overview", 14th Digital Avionics Systems Conference, November 1995.
- [5] T. K. Brown and J. A. Donaldson, "Fault Protection Design of the Command and Data Subsystem on the Cassini Spacecraft", 13th Digital Avionics Systems Conference, October 1994.
- [6] T. Kia, D.S. Bayard, and F. Tolivar, "A Precision Pointing Control System for the Space Infrared Telescope Facility," Paper AAS 97-067, 20th Annual AAS Guidance and Control Conference, Breckenridge, CO, February 5-9, 1997.
- [7] R. Jayawardhana, "Space Observatories", *Astronomy*, January 1998.
- [8] B. Pell et. al, "A Remote Agent Prototype for Spacecraft Autonomy", SPIE Proceedings Volume 2810, Denver CO, 1996.
- [9] J. Van Cleve, "SIRTf Pointing Requirements Derived from Slit Transmission in the Diffraction Limit," Report, Dept. of Astronomy, Cornell University, Ithaca, NY, May 14, 1991.